

Efficient Policy Learning

Susan Athey
athey@stanford.edu

Stefan Wager
swager@stanford.edu

Draft version February 2017

Abstract

There has been considerable interest across several fields in methods that reduce the problem of learning good treatment assignment policies to the problem of accurate policy evaluation. Given a class of candidate policies, these methods first effectively evaluate each policy individually, and then learn a policy by optimizing the estimated value function; such approaches are guaranteed to be risk-consistent whenever the policy value estimates are uniformly consistent. However, despite the wealth of proposed methods, the literature remains largely silent on questions of statistical efficiency: there are only limited results characterizing which policy evaluation strategies lead to better learned policies than others, or what the optimal policy evaluation strategies are. In this paper, we build on classical results in semiparametric efficiency theory to develop quasi-optimal methods for policy learning; in particular, we propose a class of policy value estimators that, when optimized, yield regret bounds for the learned policy that scale with the semiparametric efficient variance for policy evaluation. On a practical level, our result suggests new methods for policy learning motivated by semiparametric efficiency theory.

1 Introduction

The problem of learning treatment assignment policies is ubiquitous in applied statistics and arises, for example, in medicine when a doctor needs to decide which patients to refer to a potentially risky surgery, or in marketing when a company needs to choose which customers to send targeted offers to. A treatment assignment policy is a mapping from characteristics of units (patients or customers) to which of a set of treatments the customer should receive. Recently, new datasets have become available to researchers and practitioners that make it possible to estimate personalized policies in settings ranging from personalized offers and marketing in a digital environment to online education. In addition, technology companies, educational institutions, and researchers have begun to use explicit randomization with the goal of creating data that can be used to estimate personalized policies.

Mathematically, we formalize the problem of policy learning as follows. Suppose that we observe independent and identically distributed samples indexed by $i = 1, 2, \dots$, each characterized by features $X_i \in \mathcal{X}$, and seek a policy $\pi : \mathcal{X} \rightarrow \{\pm 1\}$ that assigns the i -th sample to receive the treatment $\pi(X_i)$, for one of two available treatments. We posit the existence of potential outcomes $Y_i(w)$, $w \in \{\pm 1\}$, that correspond to the outcome the i -th sample would experience if given treatment w (Neyman, 1923; Rubin, 1974). Given a class of policies Π , the optimal policy π^* and the regret $R(\pi)$ of any other policy are respectively

defined as

$$\pi^* = \operatorname{argmax}_{\pi \in \Pi} \{\mathbb{E}[Y_i(\pi(X_i))]\}, \quad R(\pi) = \mathbb{E}[Y_i(\pi^*(X_i))] - \mathbb{E}[Y_i(\pi(X_i))]. \quad (1)$$

Our goal is to use existing data from a randomized trial or observational study to learn a policy $\hat{\pi}$ that has low regret $R(\hat{\pi})$. Following standard practice in the statistical learning community going back to [Vapnik and Chervonenkis \(1971\)](#), our theoretical analysis will center on providing upper bounds on the decay rate of $R(\pi)$.¹

Given this setting, a natural approach to policy learning is to maximize a consistent estimate of policy value

$$Q(\pi) = \mathbb{E}[Y_i(\pi(X_i))] - \frac{1}{2}\mathbb{E}[Y_i(-1) + Y_i(+1)], \quad \hat{\pi} = \operatorname{argmax}_{\pi \in \Pi} \{\hat{Q}(\pi)\}, \quad (2)$$

where the policy value $Q(\pi)$ measures the improvement of π over a randomized baseline, and $\hat{Q}(\pi)$ is a uniformly consistent estimator for $Q(\pi)$ over the class $\pi \in \Pi$. In order to consistently estimate $Q(\pi)$, we typically need to assume that we have access to unconfounded observational data, i.e., historical data (X_i, Y_i, W_i) collected in such a way that ([Rosenbaum and Rubin, 1983](#))

$$\{Y_i(-1), Y_i(+1)\} \perp\!\!\!\perp W_i \mid X_i. \quad (3)$$

Through this paper we only assume generic regularity properties for the joint distribution of $(X_i, Y_i(-1), Y_i(+1), W_i)$. In particular, we do not assume the existence of any well-specified model; and do not assume that the class Π of policies of interest captures all sources of confounding. We discuss this statistical setup in more detail in [Section 1.1](#).

Different variants of this setup have received considerable attention across several literatures: It has been independently studied under several names, including counterfactual risk minimization ([Beygelzimer and Langford, 2009](#); [Bottou et al., 2013](#); [Dudík et al., 2011](#); [Swaminathan and Joachims, 2015](#)), outcome-weighted learning ([Chen et al., 2016](#); [Qian and Murphy, 2011](#); [Zhang et al., 2012](#); [Zhao et al., 2012](#); [Zhou et al., 2015](#)), and empirical welfare maximization ([Kitagawa and Tetenov, 2015](#)).²

Despite the wealth of literature on this topic, however, there is still no consensus on which policy value estimator $\hat{Q}(\pi)$ to use in the optimization problem (2); and, perhaps more concerning, available bounds on $R(\hat{\pi})$ are generally not sharp enough in large samples to suggest a preference for one \hat{Q} -estimator over another. As a concrete example, [Kitagawa and Tetenov \(2015\)](#) propose an estimator $\hat{\pi}$ where \hat{Q} is obtained via inverse-propensity weighting and show that, provided that the outcomes Y_i are uniformly bounded, $|Y_i| \leq M$, that the observational data we are using for learning satisfies a uniform overlap condition $\eta \leq e(x) \leq 1 - \eta$ with $e(x) = \mathbb{P}[W_i = 1 \mid X_i = x]$, and that the policy class Π is a Vapnik-Chervonenkis class with dimension $\text{VC}(\Pi)$ ([Vapnik and Chervonenkis, 1971](#)), then $\hat{\pi}$ satisfies the regret bound

$$R(\hat{\pi}) = \mathcal{O}_P \left(\frac{M}{\eta} \sqrt{\frac{\text{VC}(\Pi)}{n}} \right). \quad (4)$$

¹The reason we seek upper bounds for $R(\pi)$ —rather than exact limits—is that sharp asymptotics for discrete optimization problems of this type are often intractable; see [Bartlett and Mendelson \(2006\)](#) for a discussion.

²Estimation of optimal policies is also an important input into “contextual bandit algorithms,” multi-armed bandit algorithms that assign arriving units to different treatments based on their observable characteristics (the “context” for that individual), while balancing the tradeoff between exploration (learning about treatments whose value is uncertain) and exploitation (choosing the treatment that is best using currently available information); see, e.g., [Agarwal et al. \(2014\)](#) for a recent reference.

Now, although this bound helpfully highlights the relationship between the difficulty of policy learning and the ratio $\text{VC}(\Pi)/n$, the dependence on the data-generating distribution via the ratio M/η is too loose to speak meaningfully to our choice of \widehat{Q} -estimator. The bound (4) will hold for any reasonable choice of \widehat{Q} , e.g., ones built via matching, inverse-propensity weighting, or plug-in regression adjustments (provided the bias of $\widehat{Q}(\pi)$ at fixed values of π decays fast enough).

The goal of this paper is to provide more clarity on how to build good estimators of the form (2). We introduce some key concepts by first analyzing the problem of evaluating a single policy π as accurately as possible. Using notation from (2), we see that

$$Q(\pi) = \frac{1}{2} (\mathbb{E}[Y_i(\pi(X_i))] - \mathbb{E}[Y_i(-\pi(X_i))]), \quad (5)$$

i.e., $Q(\pi)$ is proportional to the average treatment effect in a randomized controlled trial where half the sample is assigned policy $\pi(\cdot)$, and the other half is assigned the opposite policy $-\pi(\cdot)$. Given unconfoundedness (3), there is a large literature that develops a semi-parametric efficient estimation theory for statistics like $Q(\pi)$ (Bickel et al., 1998; Hahn, 1998; Hirano et al., 2003; Robins and Rotnitzky, 1995; Robins et al., 1995), for any fixed policy π .

Our main result is that we can translate results about efficient policy evaluation into results about policy learning. Specifically, let $V(\pi)$ denote the semiparametrically efficient variance for estimating $Q(\pi)$. Furthermore, let $V_* := V(\pi^*)$ denote the semiparametrically efficient variance for evaluating the optimal policy π^* , and let V_{\max} (formally defined in (15) below) denote a sharp bound on the worst case efficient variance $\sup_{\pi} V(\pi)$ for any policy π . Then, under regularity conditions, we propose a learning rule that yields a policy $\hat{\pi}$ with regret bounded by

$$R(\hat{\pi}) = \mathcal{O}_P \left(\sqrt{V_* \log \left(\frac{V_{\max}}{V_*} \right) \frac{\text{VC}(\Pi)}{n}} \right). \quad (6)$$

We also develop regret bounds for non-parametric policy classes Π with a bounded entropy integral, such as finite-depth decision trees. Key components of our analysis include uniform concentration results for efficient policy evaluation, as well as sharp generalization bounds for weighted empirical risk minimization that may be of independent interest.

This result has several implications. First, on a conceptual level, we note that when n is large, our bound (6) is strictly better than any other regret bound for policy learning previously proposed in the literature (Beygelzimer and Langford, 2009; Kitagawa and Tetenov, 2015; Swaminathan and Joachims, 2015; Zhao et al., 2012; Zhou et al., 2015). More importantly, since our bound scales with the variance of $\widehat{Q}(\pi)$ for a single candidate policy π , this bound (6) can only be attained if $\widehat{Q}(\pi)$ is an efficient estimator of $Q(\pi)$. Thus, our bound establishes meaningful separation in terms of regret bounds available for different policy learners, and concretely establishes the relevance of the classic literature on semiparametrically efficient estimation (Bickel et al., 1998; Hahn, 1998; Hirano et al., 2003; Newey, 1994; Robins and Rotnitzky, 1995; Robins et al., 1995) to policy learning.

From a practical perspective, our result suggests novel policy learners $\hat{\pi}$. There is a recent, active, literature that seeks to use tools from optimization and machine learning to develop semiparametrically efficient average treatment effect estimation with improved finite sample performance: examples include double machine learning (Belloni et al., 2017; Chernozhukov et al., 2016), targeted learning (van der Laan and Gruber, 2010; van der Laan and Rose, 2011), and residual balancing (Athey et al., 2016a). In our experiments, we

will build policy learners that use some of these new methods to estimate $Q(\pi)$, and show that they outperform previously proposed policy learners that rely on inefficient policy estimators.

1.1 Learning Simple Policies under Semiparametric Confounding

When learning optimal policies, there are two distinct places where the features of individual units are used. First, and most obviously, an optimal policy is a mapping from features to treatments, $\pi : \mathcal{X} \rightarrow \{\pm 1\}$. In our analysis, we will consider finding the best policy within a (possibly restricted) class of policy functions Π . Second, if the data was not generated from a randomized experiment (or if the assignment probabilities vary with features in a randomized experiment), it is necessary to control for confounding features in order to attain consistent estimates of treatment effects, i.e., features that are correlated with both the treatment assignment and potential outcomes.

It is important to highlight that these two ways features may be used are distinct; therefore, throughout our analysis, the class of functions we use to control for sampling bias in our observational study is not the same as the class of functions used for policy learning. Formally, we assume that we have access to n independent and identically distributed samples $(X_i, Y_i, W_i) \in \mathcal{X} \times \mathbb{R} \times \{\pm 1\}$ generated via unconfounded potential outcomes, i.e., there are potential outcomes $\{Y_i(\pm 1)\}$ such that $Y_i = Y_i(W_i)$ and the treatment assignment satisfies the conditional independence assumption (3). We only assume generic regularity properties on this data-generating distribution, e.g., that

$$\mu_w(x) = \mathbb{E}[Y_i(w) \mid X_i = x] \text{ and } e(x) = \mathbb{E}[W_i \mid X_i = x] \quad (7)$$

belong to appropriate square-integrable Sobolev classes, and that the propensities $e(x)$ be uniformly bounded away from 0 and 1. Conversely, we will take Π to be some user-provided class with a finite entropy integral; for example, Π could be a finite family, a parametric family, or a set of tree-based functions.

Although this setup might seem surprising at first glance, we believe it to be closely in line with the classical epidemiology literature, and to be relevant in a wide variety of settings. Classical questions in causal inference using observational studies can be thought of as policy learning problems with a class of candidate policies of size $|\Pi| = 2$: For example, asking whether exercise is good for the average person is equivalent to selecting the better policy out of $\Pi = \{\pi(x) = \text{“exercise” for all } x \in \mathcal{X}, \pi(x) = \text{“no exercise” for all } x \in \mathcal{X}\}$. Thus, if we follow Rubin (1974) and seek to estimate the average causal effect of exercise while controlling for potentially non-parametric sampling bias via matching, we are exactly in our setup of interest, i.e., where π and the nuisance components $e(x)$ and $\mu_w(x)$ come from function classes of vastly different complexities.

Presently, we allow Π to be a continuum rather than a 2-point set, but the fundamental conceptual distinction between the spaces of “policy functions” and “confounding functions” persists, and is often important in real-world applications. When we want to learn a policy $\pi(\cdot)$, it is important that this policy only depend on features satisfying fairly strict desiderata: for example, they must be reliably measured and available in a deployed system, should not be manipulable by participants, and cannot discriminate based on applicable protected characteristics. Conversely, when controlling for sampling bias and/or trying to fit the moments $(\mu_w(x), e(x))$ of the data-generating function for improved efficiency, it often makes sense to control for as many pre-treatment variables as possible, especially given the recent

development of efficient methods for treatment effect estimation in high dimensions (Athey et al., 2016a; Chernozhukov et al., 2016; van der Laan and Rose, 2011).

There are of course some settings where there may be no meaningful difference between these two classes, and we are willing to learn semiparametric policy functions that are as complicated as our estimators for $(\mu_w(x), e(x))$. For example, when working with a single engineering system, e.g., a website wanting to target advertisements, we have access to a stable stream of incoming data and do not need to worry about changes in data availability or external validity. In cases like these, it makes sense to learn π using a bandit algorithm that can fit rich families of policy functions; see, e.g., Agarwal et al. (2014), Auer et al. (2002), Bastani and Bayati (2015), and references therein. However, in most public policy or epidemiology applications, we believe that using different classes of functions for $\{\pi(\cdot)\}$ versus $\{\mu_w(\cdot), e(\cdot)\}$ can be helpful, or even imperative.

1.2 Related Work

In this paper, we focus on policy learning using a single “batch” dataset collected before running any analysis (Beygelzimer and Langford, 2009; Chen et al., 2016; Bottou et al., 2013; Kallus, 2016; Kitagawa and Tetenov, 2015; Dudík et al., 2011; Swaminathan and Joachims, 2015; Qian and Murphy, 2011; Zhang et al., 2012; Zhao et al., 2012; Zhou et al., 2015). This is in contrast to the contextual bandit setup, where treatment decisions are made sequentially and so an exploration/exploitation trade-off arises (e.g., Agarwal et al., 2014; Auer et al., 2002; Bastani and Bayati, 2015). In the batch setting, all exploration happened in the past and the future is only about exploitation, so the policy learning problem reduces to a statistical problem of extracting as much information as possible from already collected data.

The concrete algorithms we propose for policy learning—see Section 2.1 for details—most closely resemble methods developed by Dudík et al. (2011) and Zhang et al. (2012), who use doubly robust estimates for $\hat{Q}(\pi)$ that are well-known to be semiparametrically efficient under appropriate conditions (Hahn, 1998; Robins and Rotnitzky, 1995). The above papers, however, did not study the potential for efficiency gains from this method, and only focused on doubly robust consistency. We also note the recent work of Zhou et al. (2015), who uses a form of residualization to improve on outcome-weighted learning as proposed by Zhao et al. (2012); however, the scoring method they use is inefficient whenever treatment propensities may deviate from $e(x) = 1/2$, and their theoretical analysis is not sharp enough to provide regret bounds that scale with the second moment of the scores.

Finally, from a theoretical perspective, there has been considerable interest in concentration bounds for weighted learning that, in our setting, could be used to obtain regret bounds of the type (Cortes et al., 2010; Maurer and Pontil, 2009; Swaminathan and Joachims, 2015)

$$R(\hat{\pi}) = \mathcal{O}_P \left(\sqrt{V_* \frac{\log(n) \text{VC}(\Pi)}{n}} \right). \quad (8)$$

The bound in (8) has the desirable property that the variance of the policy value estimator $\hat{Q}(\pi^*)$ enters the bound. However, relative to the bounds obtained in this paper, (8) has an extraneous $\log(n)$ factor, so that the bound is strictly larger than classical bounds of the form (4) when n is large. A contribution of our analysis is to tighten the bound to eliminate the $\log(n)$ term. We also note that Swaminathan and Joachims (2015) explicitly consider the problem of policy learning using “empirical Bernstein” bounds of the above type for

variance penalization; however, they use inverse-propensity weighting rather than efficient estimators for $\widehat{Q}(\pi)$.

2 From Efficient Policy Evaluation to Learning

Recall that we are interested in the following problem: We have n i.i.d. samples (X_i, Y_i, W_i) drawn from a regular, unconfounded distribution as described in (7). Given such a data-generating distribution, we want to learn a policy assignment rule $\hat{\pi} : \mathcal{X} \rightarrow \{\pm 1\}$ belonging to a class $\hat{\pi} \in \Pi$, such as to make the regret (1) small. Following an extensive existing literature (Beygelzimer and Langford, 2009; Bottou et al., 2013; Chen et al., 2016; Dudík et al., 2011; Kitagawa and Tetenov, 2015; Swaminathan and Joachims, 2015; Zhao et al., 2012; Zhou et al., 2015), we focus on learners $\hat{\pi}$ obtained by optimizing a policy value estimate $\widehat{Q}(\pi)$ as in (2). Our goal is to find a class of efficient \widehat{Q} -estimators that yield $\hat{\pi}$ -learners who inherit their efficiency properties.

Perhaps the simplest way to construct semiparametrically efficient estimators for $Q(\pi)$ is via “double machine learning” (Belloni et al., 2017; Chernozhukov et al., 2016), i.e., estimators of the following form. We first divide the data into K evenly-sized folds and, for each fold $k = 1, \dots, K$, run a machine learning estimator of our choice on the other $K - 1$ data folds to estimate the functions $\mu_{\pm 1}(x)$ and $e(x)$; denote the resulting estimates $\mu_{\pm 1}^{(-k)}(x)$ and $\hat{e}^{(-k)}(x)$. Then, given these pre-computed values, we estimate $Q(\pi)$ as

$$\widehat{Q}_{DML}(\pi) = \frac{1}{n} \sum_{i=1}^n \pi(X_i) \widehat{\Gamma}_i, \quad \widehat{\Gamma}_i := \hat{\mu}_{+1}^{(-k(i))}(X_i) - \hat{\mu}_{-1}^{(-k(i))}(X_i) + W_i \frac{Y_i - \hat{\mu}_{W_i}^{(-k(i))}(X_i)}{\hat{e}_{W_i}^{(-k(i))}(X_i)}, \quad (9)$$

where $k(i) \in \{1, \dots, K\}$ denotes the fold containing the i -th conservation. Here, we have also used the short-hand

$$\hat{e}_{W_i}^{(-k(i))}(X_i) = \frac{1}{2} + W_i \left(\frac{1}{2} - \hat{e}^{(-k(i))}(X_i) \right) \quad (10)$$

to denote estimates of the class-specific propensities. The K -fold algorithmic structure used in (9) was proposed by Schick (1986) as a general purpose tool for efficient estimation in semiparametric models, and was independently called “cross-fitting” by Chernozhukov et al. (2016) and “cross-estimation” by Wager et al. (2016).

Under weak assumptions, Belloni et al. (2017) and Chernozhukov et al. (2016) show that the double machine learning estimator (9) achieves the semiparametrically efficient rate for estimating $Q(\pi)$ (Hirano et al., 2003),

$$\sqrt{n} \left(\widehat{Q}_{DML}(\pi) - Q(\pi) \right) \Rightarrow \mathcal{N}(0, V(\pi)), \quad (11)$$

$$V(\pi) = \text{Var} [\pi(X)\tau(X)] + \mathbb{E} \left[\frac{\text{Var} [Y(-1) | X = X_i]}{1 - e(X_i)} + \frac{\text{Var} [Y(+1) | X = X_i]}{e(X_i)} \right], \quad (12)$$

provided the root-mean squared error of the estimators $\hat{\mu}_{\pm 1}(x)$ and $\hat{e}(x)$ goes to zero faster than $1/n^{1/4}$. For a review of conditions under which such convergence is possible, see Chernozhukov et al. (2016); for example, it is enough that $e(\cdot)$ and $\mu_w(\cdot)$ belong to appropriate L_2 -Sobolev classes. This estimator is closely related to the classical semiparametric two-stage methods studied by, e.g., Hahn (1998), Newey (1994) and Robins and Rotnitzky

(1995). A contribution of this paper is to extend the theoretical results from this literature to the policy learning problem, where we need uniform convergence of estimates of $Q(\pi)$ for $\pi \in \Pi$ in order to translate results about policy evaluation into regret bounds for the learned policies.

One of our main results is that by optimizing the value estimate obtained via double machine learning, we attain strong regret bounds for policy learning. Throughout our analysis, we will make the following assumption about the machine learning method underlying (9).

Assumption 1 (Consistent machine learning). Whenever we use a double machine learning estimator \hat{Q}_{DML} constructed as in (9), we assume that the machine learning methods used to construct our estimator satisfy the following consistency guarantees. The methods must be uniformly consistent,

$$\sup_{x \in \mathcal{X}} |\hat{\mu}_w(x) - \mu_w(x)|, \quad \sup_{x \in \mathcal{X}} |\hat{e}_{\pm 1}(x) - e_{\pm 1}(x)| \rightarrow_p 0, \quad (13)$$

and, moreover, the methods must converge as $n^{1/4}$ under L_2 loss,

$$\mathbb{E} \left[(\hat{\mu}_w(X) - \mu_w(X))^2 \right]^{\frac{1}{2}}, \quad \mathbb{E} \left[\left(\frac{1}{\hat{e}_{\pm 1}(X)} - \frac{1}{e_{\pm 1}(X)} \right)^2 \right]^{\frac{1}{2}} \leq \frac{\sqrt{a(n)}}{n^{1/4}}, \quad (14)$$

for some sequence $a(n) \rightarrow 0$, where X is taken to be an independent test example drawn from the same distribution as the training data.

To get a feeling for the types of results we can obtain for policy learning with double machine learning, we consider below the case where Π belongs to a VC class. This setup allows us to state a result without resorting to too much notation. As discussed in the introduction, our results will also depend on the worst-case efficient variance V_{\max} for estimating any policy. Given our data-generating assumptions, we have a simple form for V_{\max} , namely

$$V_{\max} = \mathbb{E} [\tau^2(X)] + \mathbb{E} \left[\frac{\text{Var} [Y(-1) | X = X_i]}{1 - e(X_i)} + \frac{\text{Var} [Y(+1) | X = X_i]}{e(X_i)} \right], \quad (15)$$

and note that $V(\pi) = V_{\max} - Q^2(\pi)$ for any policy π .

Theorem 1. Define $\hat{Q}_{DML}(\pi)$ as in (9), and let $\hat{\pi} = \text{argmin}_{\pi \in \Pi} \hat{Q}_{DML}(\pi)$. Given Assumption 1, suppose moreover that we have overlap, i.e.,

$$\eta \leq \mathbb{P} [W = 1 | X = x] \leq 1 - \eta \text{ for some } \eta > 0, \quad (16)$$

that the irreducible noise $\varepsilon_i = Y_i - \mathbb{E} [Y_i | X_i, W_i]$ is uniformly sub-Gaussian conditionally on X_i and W_i , and that Π is a VC class of dimension $\text{VC}(\Pi)$. Then, for any $\delta > 0$, there is a universal constant³ C_δ , as well as a threshold N that depends on the constants used to define the regularity assumptions such that, with probability at least $1 - \delta$,

$$R(\hat{\pi}) \leq C_\delta \sqrt{\text{VC}(\Pi) V(\pi^*) \log \left(\frac{V_{\max}}{V(\pi^*)} \right)} / n, \text{ for all } n \geq N, \quad (17)$$

where $V(\pi)$ denotes the semiparametric efficient variance for policy evaluation (12) and V_{\max} is as defined in (15).

³Throughout this paper, we will use C and C_δ to denote different universal constants; no two instantiations of C and C_δ should be assumed to denote the same constant.

We develop the technical tools required to prove this result in Section 3; before that, however, we discuss implementation of the method and present a simple example.

2.1 Implementation via Weighted Classification

In the previous sections, we established regret bounds for the policy $\hat{\pi}_{DML}$ obtained by maximizing $\hat{Q}_{DML}(\pi)$ is as defined in (9). In order to carry out this optimization, we follow [Beygelzimer and Langford \(2009\)](#), [Kitagawa and Tetenov \(2015\)](#), [Zhang et al. \(2012\)](#), [Zhao et al. \(2012\)](#) and [Zhou et al. \(2015\)](#), and note that $\hat{\pi}_{DML}$ can also be understood as the empirical risk minimizer in a weighted classification problem:

$$\hat{\pi} = \operatorname{argmax}_{\pi \in \Pi} \left\{ \frac{1}{n} \sum_{i=1}^n \lambda_i Z_i \pi(X_i) \right\}, \quad \lambda_i = |\hat{\Gamma}_i|, \quad Z_i = \operatorname{sign}(\hat{\Gamma}_i), \quad (18)$$

i.e., we want to train a classifier with response Z_i using weights λ_i ; recall that $\hat{\Gamma}_i$ was defined in (9). Given this formulation as a weighted classification problem, we can use standard off-the-shelf tools for weighted classification to learn $\hat{\pi}$, e.g., classification trees ([Breiman et al., 1984](#)) or support-vector machines ([Cortes and Vapnik, 1995](#)).⁴

Several other proposals also fit into this framework, with different choices of $\hat{\Gamma}_i$. [Zhao et al. \(2012\)](#) assume a randomized controlled trial and use $\hat{\Gamma}_i = W_i Y_i / \mathbb{P}[W_i = 1]$, while [Kitagawa and Tetenov \(2015\)](#) use inverse-propensity weighting $\hat{\Gamma}_i = W_i Y_i / \hat{e}_{W_i}(X_i)$. In an attempt to stabilize the weights, [Beygelzimer and Langford \(2009\)](#) introduce an “offset”

$$\hat{\Gamma}_i = \frac{W_i}{\hat{e}_{W_i}(X_i)} \left(Y_i - \frac{\max\{Y_i\} + \min\{Y_i\}}{2} \right),$$

while [Zhou et al. \(2015\)](#) go further and advocate

$$\hat{\Gamma}_i = \frac{W_i}{\hat{e}_{W_i}(X_i)} \left(Y_i - \frac{\hat{\mu}_{+1}(X_i) + \hat{\mu}_{-1}(X_i)}{2} \right).$$

None of the above methods, however, are built on semiparametrically efficient policy evaluation, and so they do not fit into the class of algorithms covered by Theorem 1. Finally, the method advocated by [Zhang et al. \(2012\)](#) actually takes the same form as our procedure (18). However, the paper by [Zhang et al. \(2012\)](#) does not provide regret bounds; moreover, they do not use “cross-fitting” or “cross-estimation” as in (9), so it is unclear under what conditions their method satisfies the bounds from Theorem 1.⁵

2.2 A Simple Illustration

We illustrate our approach with a simple simulation example. Suppose that we have access to data $(X, Y, W) \in [-1, 1]^2 \times \mathbb{R} \times \{\pm 1\}$, and want to learn a policy function $\pi(\cdot)$ that

⁴In our discussion so far, we have largely left aside questions on how to estimate $\hat{\mu}_w(\cdot)$ and $\hat{e}(\cdot)$, provided that they are obtained using some machine learning method that satisfies Assumption 1. However, in terms of finite-sample performance, experience suggests that it may be preferable to use methods for $\hat{\mu}_{+1}(\cdot)$ and $\hat{\mu}_{-1}(\cdot)$ that estimate both arms simultaneously while explicitly seeking out treatment effect heterogeneity ([Athey et al., 2016b](#); [Imai and Ratkovic, 2013](#)), and to use methods for $\hat{e}(\cdot)$ that privilege balance of inverse-propensity weighted features over classification accuracy ([Athey et al., 2016a](#); [Graham et al., 2012](#)).

⁵Other related methods, such as those by [Dudík et al. \(2011\)](#) and [Swaminathan and Joachims \(2015\)](#), apply in a setting with multiple available treatments and so do not directly fit into the setting of (18). Extending our efficiency analysis to the multi-treatment regime would be of considerable interest.




Figure 1: Results from two attempts at learning a policy π by counterfactual risk minimization (2) over depth-2 decision trees, as described in Section 2.2, with \hat{Q} -estimators obtained by both inverse-propensity weighting and double machine learning. The dashed blue line denotes the optimal decision rule (treat in the upper-right corner, do not treat elsewhere); the solid black lines denote learned policies π (treat in the shaded regions, do not treat elsewhere). We also use a heat map to depict the average policies learned across 200 simulations; the darkest red regions are never treated and lightest yellow regions are always treated. Simulations were performed using $n = 500$ samples drawn according to (19).

can be written as a depth-2 decision tree. Here, the data is independent and identically distributed as

$$X_i \sim U([-1, 1]^2), W_i | X_i \sim \text{Bern}(e(x)), Y_i | W_i, X_i \sim \mathcal{N}(\mu(X_i) + \tau(X_i)W_i/2, 1), \quad (19)$$

with $e(x) = 1/(1 + e^{-(x_1 + x_2)})$, $\mu(x) = 2e^{-(x_1 + x_2)}$, and $\tau(x) = 2/[(1 + e^{-4x_1})(1 + e^{-4x_2})] - 0.4$. For our purposes, the salient facts about this data-generating distribution are that it is unconfounded (3); however, it cannot be represented by trees, and the optimal policy is not in our class Π of interest, i.e., depth-2 decision trees.

In Figure 1, we show results for learning π using two different counterfactual risk minimization strategies of the form (2), but with different \hat{Q} -estimators. The left panel obtains \hat{Q} by inverse-propensity weighting; conversely, the right panel uses an efficient double machine learning \hat{Q} -estimator. We see that the policies $\hat{\pi}$ learned via efficient policy evaluation are much better than those learned by inverse-propensity weighting. Across 200 simulation runs, inverse-propensity weighting led to a mean regret of 0.143, whereas double machine learning got a mean regret of 0.063 relative to the best possible depth-2 tree. Figure 1 shows both a single realization of each method, as well as the average policy learned across 200 simulation runs.

In terms of specifics, we started by learning $\hat{\mu}_{\pm 1}(\cdot)$ and $\hat{e}(\cdot)$ via a lasso (Tibshirani, 1996) on a polynomial basis expansion (with interactions), all while using out-of-fold prediction as in (9). For $\hat{\mu}_{\pm 1}(\cdot)$, we fit both response functions simultaneously while writing them in terms of a main effect $(\hat{\mu}_{+1}(\cdot) + \hat{\mu}_{-1}(\cdot))/2$ and a treatment effect $(\hat{\mu}_{+1}(\cdot) - \hat{\mu}_{-1}(\cdot))/2$; when

the treatment effect is weaker than the main effect, this re-parametrization can interact with the lasso penalty in a way that improves the resulting fit (Imai and Ratkovic, 2013). Finally, we got $\hat{\pi}$ by optimizing (18) over depth-2 trees with the R-package `evtree`, which learns an optimal classification tree using an evolutionary algorithm (Grubinger et al., 2014). Inverse-propensity weighting uses $\hat{\Gamma}_i = Y_i/\hat{e}_{W_i}(X_i)$, while our method uses the form in (9).

3 Theoretical Development

As is common in the literature on semiparametric estimation, our proof is built around a study of the efficient influence function for policy evaluation,

$$\tilde{Q}(\pi) = \frac{1}{n} \sum_{i=1}^n \pi(X_i) \Gamma_i, \quad \Gamma_i := \mu_{+1}(X_i) - \mu_{-1}(X_i) + W_i \frac{Y_i - \mu_{W_i}(X_i)}{e_{W_i}(X_i)}. \quad (20)$$

Note that $\tilde{Q}(\pi)$ can be understood as a version of $\hat{Q}(\pi)$ computed by an oracle who has access to the true functions $\mu_w(x)$ and $e(x)$. Many classical results in semiparametric efficiency theory rely on characterizing the realizable estimators $\hat{Q}(\pi)$ in terms of the oracle quantities $\tilde{Q}(\pi)$: Classical results going back at least to Newey (1994) show that, under appropriate regularity conditions, two-stage estimators of the form (9) are asymptotically equivalent to an average of the efficient influence function,

$$\sqrt{n} \left(\hat{Q}(\pi) - \tilde{Q}(\pi) \right) \rightarrow_p 0. \quad (21)$$

This type of idea underlies, among other things, the analysis used to justify double machine learning (Belloni et al., 2017; Chernozhukov et al., 2016).

Our proof is structured as follows. Having spelled out assumptions about the policy class below, we proceed in Section 3.2 to prove a regret bound that would be available to a statistician who could optimize the infeasible objective $\tilde{Q}(\pi)$ rather than the feasible double machine learning objective $\hat{Q}(\pi)$. Then in Section 3.3, we establish a strengthening of (21) that holds uniformly over all $\pi \in \Pi$, and use this coupling to establish a first result about learning with $\hat{Q}(\pi)$. Finally, in Section 3.4, we re-visit and strengthen our bounds under stronger assumptions on Π that hold, for example, when Π is a VC-class.

3.1 Assumptions about the Policy Class

Although we stated our first result, Theorem 1, under the simple assumption that Π be a Vapnik-Chervonenkis class, we will develop our technical results under weaker, more abstract assumptions on Π . In order to obtain regret bounds as in (17) that decay as $1/\sqrt{n}$, we of course need some control over the complexity of the class Π .

Here, we do so using bounds on the Hamming entropy of Π . For any discrete set of points $\{X_1, \dots, X_m\}$ and any $\varepsilon > 0$, define the ε -Hamming covering number $N_H(\varepsilon, \Pi, \{X_1, \dots, X_m\})$ as the smallest number of policies $\pi : \{X_1, \dots, X_m\} \rightarrow \{\pm 1\}$ required to ε -cover Π under Hamming distance,

$$H(\pi_1, \pi_2) = \frac{1}{m} \sum_{j=1}^m \mathbf{1}(\{\pi_1(X_j) \neq \pi_2(X_j)\}).$$

Then, define the ε -Hamming entropy of Π as $\log(N_H(\varepsilon, \Pi))$, where

$$N_H(\varepsilon, \Pi) = \sup \{N_H(\varepsilon, \Pi, \{X_1, \dots, X_m\}) : X_1, \dots, X_m \in \mathcal{X}; m \geq 1\}$$

is the largest possible number of functions needed the ε -cover Π under Hamming distance for any discrete set of points. We note that this notion of entropy is purely geometric, and does not depend on the distribution used to generate the X_i . Finally, we assume the following:

Assumption 2 (Entropy bound). We assume that there is a constant $C, \omega > 0$ such that the Hamming entropy of Π is bounded by $\log(N_H(\varepsilon, \Pi)) \leq C\varepsilon^{\omega-1}$ for all $0 < \varepsilon < 1$.

Given this assumption define the complexity κ of the class Π in terms of a variant of the classical entropy integral of [Dudley \(1967\)](#):⁶

$$\kappa(\Pi) = \int_0^1 \sqrt{\log(N_H(\varepsilon^2, \Pi))} d\varepsilon. \quad (22)$$

Using entropy integrals to bound model class complexity is ubiquitous in empirical process theory (see, e.g., [Boucheron et al. \(2013\)](#)), and easily allow us to specialize to more restrictive cases. In particular, if Π is a Vapnik-Chervonenkis class, [Haussler \(1995\)](#) showed that

$$\log(N_H(\varepsilon, \Pi)) \leq d(\log(\varepsilon^{-1}) + \log(2) + 1) + \log(d + 1) + 1, \text{ with } d := \text{VC}(\Pi), \quad (23)$$

from which we can check that $\kappa^2(\Pi) \leq 6d$ for any value of $d = 1, 2, \dots$. Meanwhile, a quick analysis of depth- L decision trees with $X_i \in \mathbb{R}^d$ can establish that $\log(N_H(\varepsilon, \Pi)) = \mathcal{O}(2^L \log(\varepsilon^{-1}) + 2^L \log(d) + L2^L)$, and so Assumption 2 again holds.⁷

3.2 Rademacher Complexities and Oracle Regret Bounds

We start our analysis by characterizing the regret of an oracle learner who has access to bound $\mu_{\pm}(\cdot)$ and $e(\cdot)$, and chooses their policy $\hat{\pi}$ by optimizing the infeasible value estimator $\tilde{Q}(\pi)$ as defined in (20). The advantage of studying this oracle is that it allows us, for the time being, to abstract away from the specific machine learning methods used to obtain $\tilde{Q}(\pi)$, and instead to focus on the complexity of counterfactual risk minimization over the class Π .

Specifically, our present goal is to study concentration of the empirical process $\tilde{Q}(\pi) - Q(\pi)$ for all $\pi \in \Pi$. Recalling the definition of Γ_i from (20), a convenient way to bound the supremum of our empirical process of interest is by controlling its Rademacher complexity $\mathcal{R}(\Pi)$, defined as

$$\mathcal{R}(\Pi) = \sup_{\pi \in \Pi} \left\{ \frac{1}{n} \sum_{i=1}^n Z_i \Gamma_i \pi(X_i) \right\} \quad (24)$$

where the Z_i are independent Rademacher (i.e., sign) random variables $Z_i = \pm 1$ with probability 1/2 each ([Bartlett and Mendelson, 2002](#)). For intuition as to why Rademacher complexity is a natural complexity measure, note that $\mathcal{R}(\Pi)$ characterizes the maximum (weighted) in-sample classification accuracy on randomly generated labels Z_i over classifiers $\pi \in \Pi$; thus, $\mathcal{R}(\Pi)$ directly measures how much we can overfit to random coin flips using Π .

⁶Assumption 2 immediately guarantees that this integral is finite.

⁷To verify this result for trees, we follow [Bartlett and Mendelson \(2002\)](#) and view each tree-leaf as a conjunction of L boolean functions, along with a sign. A simple argument then shows that a library of $4d^2 L 2^L \varepsilon^{-1}$ boolean functions lets us approximate each leaf to within Hamming error $2^{-L}\varepsilon$; and so we can also approximate the tree to within ε Hamming error. The resulting bound on $N_H(\varepsilon, \Pi)$ follows by noting that a full tree has $2^L - 1$ splits, and so can be approximated using $2^L - 1$ of these boolean functions.

Following this proof strategy, we start by providing a bound for $\mathcal{R}(\Pi)$ below that scales as $\sqrt{\mathbb{E}[\Gamma^2]}/\sqrt{n}$. Despite its simple form, we are not aware of existing proofs of such results in the literature. Bounds that scale as $\max\{\Gamma_i\}/\sqrt{n}$ are standard, but are not strong enough to move past results of the type (4) in our setting. Meanwhile, bounds that scale as $\sqrt{\mathbb{E}[\Gamma^2] \log(n)/n}$ are developed by [Cortes et al. \(2010\)](#) and [Maurer and Pontil \(2009\)](#); however, the additional $\log(n)$ factor makes these bounds inappropriate for asymptotic analysis.

Lemma 2. *Suppose that the class Π satisfies Assumption 2, and that the weights Γ_i in (24) are drawn from a sub-Gaussian distribution. Then, there is a universal constant $C > 0$ for which*

$$\mathbb{E}[\mathcal{R}(\Pi)] \leq 8(\kappa(\Pi) + C) \sqrt{\frac{\mathbb{E}[\Gamma_i^2]}{n}} + \mathcal{O}\left(\frac{\log(n)}{n}\right), \quad (25)$$

where $\kappa(\Pi)$ is the complexity of Π as defined in (22).

Given this Rademacher complexity bound, we can obtain a uniform concentration bound for $\tilde{Q}(\pi)$ using standard arguments. Here, we refine an argument of [Bartlett and Mendelson \(2002\)](#) using Talagrand's inequality to obtain a bound that scales as $\sqrt{\mathbb{E}[\Gamma_i^2]}$ rather than $\sup|\Gamma_i|$. In the statement of the result below, we note that $V_{\max} = \mathbb{E}[\Gamma_i^2]$, as is clear from (15).

Theorem 3. *Under the conditions of Lemma 2, the averaged efficient influence functions $\tilde{Q}(\pi)$ concentrate uniformly: There is a universal constant $C > 0$ such that, for any $\delta > 0$, with probability at least $1 - \delta$,*

$$\sup_{\pi \in \Pi} \left| \tilde{Q}(\pi) - Q(\pi) \right| \leq \left(16\kappa(\Pi) + \sqrt{2 \log(\delta^{-1})} + C \right) \sqrt{\frac{V_{\max}}{n}} + \mathcal{O}\left(\frac{\log(n)}{n}\right). \quad (26)$$

In particular, if we set $\tilde{\pi} \in \operatorname{argmax}\{\tilde{Q}(\pi) : \pi \in \Pi\}$, then again with probability at least $1 - \delta$,

$$R(\tilde{\pi}) \leq 2 \left(16\kappa(\Pi) + \sqrt{2 \log(\delta^{-1})} + C \right) \sqrt{\frac{V_{\max}}{n}} + \mathcal{O}\left(\frac{\log(n)}{n}\right), \quad (27)$$

where $R(\tilde{\pi})$ stands for the regret of policy $\tilde{\pi}$.

3.3 Uniform Coupling with the Efficient Score

In the previous section, we established risk bounds that would hold if we could optimize the infeasible value function $\tilde{Q}(\pi)$; we next need to extend these bounds to cover the situation where we optimize a feasible value function. As discussed above, we focus on the double machine learning estimator $\hat{Q}_{DML}(\pi)$. For a single, fixed policy π , [Chernozhukov et al. \(2016\)](#) showed that $\tilde{Q}(\pi)$ and $\hat{Q}_{DML}(\pi)$ are asymptotically equivalent, meaning that the discrepancy between the two value estimates decays faster than the variance of either. However, in our setting, the statistician gets to optimize over all policies $\pi \in \Pi$, and so coupling results established for a single pre-determined policy π are not strong enough.

The following lemma extends the work of [Chernozhukov et al. \(2016\)](#) to the case where we seek to establish a coupling of the form (21) that holds simultaneously for all $\pi \in \Pi$. We establish a perhaps surprising strong result: The bound (28) is the same bound as we might expect to obtain for a single policy π as in (21), and the complexity $\kappa(\Pi)$ of the class Π does not affect the leading-order constants in the bound.

Lemma 4. *Under the conditions of Lemma 2, suppose that we obtain $\widehat{Q}_{DML}(\pi)$ by double machine learning according to Assumption 1. Then, for any $\delta > 0$, there is a universal constant $C_\delta \geq 0$ for which*

$$\begin{aligned} & \sup \left\{ \left| \widehat{Q}_{DML}(\pi) - \widetilde{Q}(\pi) \right| : \pi \in \Pi \right\} \\ & \leq C_\delta a \left(\left\lfloor \frac{K-1}{K} n \right\rfloor \right) / \sqrt{\left\lfloor \frac{K-1}{K} n \right\rfloor} + \mathcal{O}_P \left(\frac{1}{n^{3/4}} \right), \end{aligned} \quad (28)$$

where K is the number of data folds used for double machine learning (9). The remainder term \mathcal{O}_P may depend on problem-specific parameters, such as $\kappa(\Pi)$.

The following result is an immediate corollary of Theorem 3 and Lemma 4. In terms of assumptions, we note that assuming overlap as in (16) and sub-Gaussianity of the irreducible noise $\varepsilon_i = Y_i - \mathbb{E}[Y_i | X_i, W_i]$ lets us guarantee that the weights Γ_i in (24) are sub-Gaussian.

Theorem 5. *Suppose that Assumption 1 and 2 hold, that we have overlap as in (16), and that the irreducible noise $\varepsilon_i = Y_i - \mathbb{E}[Y_i | X_i, W_i]$ is uniformly sub-Gaussian conditionally on X_i and W_i . Then, for any $\delta > 0$, there is a universal constant C_δ and a problem-specific threshold N such that*

$$R(\hat{\pi}) \leq C_\delta \max \{ \kappa(\Pi), 1 \} \sqrt{\frac{V_{\max}}{n}} \quad (29)$$

with probability at least $1 - \delta$ for all $n \geq N$, where $\hat{\pi}$ optimizes the double machine learning risk estimate as in (2).

The above bound is close to our desideratum: It obtains regret bounds for a realizable policy that have the desired dependence on $\kappa(\Pi)$ and n (recall that, for VC-classes, $\kappa(\Pi) = \mathcal{O}(\sqrt{\text{VC}(\Pi)})$), and scales with a semiparametric variance bound rather than with a crude bound on, e.g., $\sup |Y_i|$. However, a down-side of the bound (29) is that it depends on V_{\max} , i.e., a worst-case bound for the semiparametric efficient variance for evaluating any policy, and not as V_* , the semiparametric efficient variance for evaluating the optimal policy π^* . In the following section, we seek to replace the dependence on V_{\max} with one on V_* , at the expense of stronger assumptions on the class Π .

3.4 Improved Bounds via Slicing

To move past the V_{\max} scaling above, we need a closer analysis that cuts the space Π into strata of policies that have comparable values of $V(\pi)$, and then develop concentration bounds for these strata separately. This “slicing” idea is common in the literature, and has been used in different contexts by, e.g., Bartlett et al. (2005) and Giné and Koltchinskii (2006).

The reason we might expect slicing to work in our case is that, as discussed earlier, the efficient variance for evaluating any given policy π is $V(\pi) = V_{\max} - Q^2(\pi)$. Thus, any “good” policy, i.e., with a large value $Q(\pi)$, must also have a small efficient variance $V(\pi)$. More specifically, letting Π_λ denote the set of policies with with regret at most λ ,

$$\Pi_\lambda = \{ \pi \in \Pi : R(\pi) \leq \lambda \}. \quad (30)$$

we immediately see that

$$\sup \{ V(\pi) : \pi \in \Pi_\lambda \} \geq V_\lambda := V(\pi^*) + 2\lambda Q(\pi^*), \quad (31)$$

where π^* is an optimal policy.

This improved variance bound suggests an argument proceeding in two stages: First, Theorem 5 already established that the learned policy $\hat{\pi}$ has regret going to 0 and so $\mathbb{P}[\hat{\pi} \in \Pi_\lambda] \rightarrow 1$ for any fixed $\lambda > 0$; then, in a second stage, we use the improved variance bounds in (31) to get tighter concentration bounds for $\hat{\pi}$.

The fact that such a slicing argument works is not to be taken for granted; and, in our case, is a property that hinges crucially on the fact that, if we use an efficient method for policy evaluation, then there do not exist any policies π that simultaneously have low regret and are hard to evaluate (i.e., $V(\pi)$ is large). Given other evaluation methods, e.g., inverse propensity weighting as used in Swaminathan and Joachims (2015), such a slicing argument may not work; and in fact Swaminathan and Joachims (2015) need to add a form of auxiliary variance penalization to their estimator in order to obtain regret bounds that scale with V_* rather than V_{\max} .

Below we establish a concentration bound for Π_λ under the following assumption on the entropy of Π : for some $\alpha > 0$,

$$\log(N_H(\varepsilon, \Pi)) \leq \alpha \log(\varepsilon^{-1}), \text{ for all } 0 < \varepsilon < \frac{1}{2}. \quad (32)$$

Note that if Π is a VC-class then the above holds with $\alpha \leq 6 \text{VC}(\Pi)$; see (23).

Theorem 6. *Under the conditions of Theorem 5, suppose moreover that (32) holds and let $\lambda > 0$ be predetermined. Then, for any $\delta > 0$ there is a universal constant C , as well as a potentially problem-specific threshold N , such that for all $n \geq N$, with probability at least $1 - \delta$,*

$$\sup_{\pi \in \Pi_\lambda} \left| \tilde{Q}(\pi) - Q(\pi) \right| \leq C_\delta \sqrt{\frac{\alpha}{n} V_\lambda \log\left(\frac{V_{\max}}{V_\lambda}\right)}, \quad (33)$$

where α controls the complexity of Π via (32).

Although the above bound may superficially look like a direct extrapolation from (26), we caution that the proof relies on a subtly different construction than that used in Lemma 2, requiring stronger assumptions. In particular, the additional factor $\log(V_{\max}/V_*)$ in the bound below is directly tied to the entropy growth rate assumed in (32); and in fact is closely related to the $\log(n)$ factor appearing in the bounds of Cortes et al. (2010) and Maurer and Pontil (2009).

Given this result, the proof of our main result follows immediately. Note that, whenever the condition (32) holds, we can replace the term $\text{VC}(\Pi)$ with $\alpha/6$ in the statement of Theorem 1.

Proof of Theorem 1. Given that Π is a VC class, recall that (32) holds with $\alpha \leq 6 \text{VC}(\Pi)$. Now, set $\lambda = V(\pi^*)/(2Q(\pi^*))$, so $V_\lambda = 2V(\pi^*)$. By Theorem 5, we know that $\mathbb{P}[\hat{\pi} \in \Pi_\lambda] \rightarrow 1$, and so (17) follows immediately from Theorem 6 paired with Lemma 4. \square

4 Discussion

In this paper, we showed how classical concepts from the literature on semiparametric efficiency can be used to develop performant algorithms for policy learning with strong asymptotic guarantees. Our regret bounds may prove to be particularly relevant in applications

since, unlike existing bounds, they are sharp enough to distinguish between different a priori reasonable policy learning schemes (e.g., ones based on inverse-propensity weighting versus double machine learning), and thus provide methodological guidance to practitioners.

More generally, our experience shows that results on semiparametrically efficient estimation are not just useful for statistical inference, but are also directly relevant to applied decision making problems. It will be interesting to see whether related insights will prove to be more broadly helpful for, e.g., sequential problems with contextual bandits, or non-discrete decision making problems involving, say, price setting or capacity allocation.

5 Proofs

5.1 Proof of Lemma 2

Our proof of this result follows the outline of Dudley’s chaining argument, whereby we construct a sequence of approximating sets of increasing precision for $\tilde{Q}(\pi)$ with $\pi \in \Pi$, and then use finite concentration inequalities to establish describe the behavior of $\tilde{Q}(\pi)$ over this approximation set. The improvements in our results relative to existing bounds described in the body of the text come from the details of the proof, that is, a careful construction of approximating sets targeted to the problem at hand, and the use of sharp concentration inequalities.

Given these preliminaries, we start by defining the conditional 2-norm distance between two elements $\pi_1, \pi_2 \in \Pi$ as

$$D^2(\pi_1, \pi_2) = \frac{1}{4} \sum_{i=1}^n \Gamma_i^2 (\pi_1(X_i) - \pi_2(X_i))^2 / \sum_{i=1}^n \Gamma_i^2, \quad (34)$$

and let $N_D(\varepsilon, \Pi, \{X_i \Gamma_i\})$ be the ε -covering number in this distance. To bound N_D , imagine creating another sample $\{X'_j\}_{j=1}^m$, with X'_j contained in the support of $\{X_i\}_{i=1}^n$, such that

$$\left| \left| \{j \in 1, \dots, m : X'_j = X_i\} \right| - m \Gamma_i^2 / \sum_{i=1}^n \Gamma_i^2 \right| \leq 1.$$

We immediately see that, for any two policies π_1 and π_2 ,

$$\frac{1}{m} \sum_{j=1}^m 1(\{\pi_1(X'_j) \neq \pi_2(X'_j)\}) = D^2(\pi_1, \pi_2) + \mathcal{O}\left(\frac{1}{m}\right).$$

Now, recall that N_H as used our entropy integral (22) does not depend on sample size, so we can without reservations make m arbitrarily large, and conclude that

$$N_D(\varepsilon, \Pi, \{X_i, \Gamma_i\}) \leq N_H(\varepsilon^2, \Pi). \quad (35)$$

In other words, we have found that we can bound the D -entropy of Π with respect to its distribution-independent Hamming entropy.

Now, for every element $\pi \in \Pi$, define a set of approximations $A_j(\pi) : \mathcal{X} \rightarrow \{\pm 1\}$ for $j = 1, 2, \dots$ with the property that $D(A_j(\pi), A_{j+1}(\pi)) \leq 2^{-j}$, and that the set $\{A_j(\pi) : \pi \in \Pi\}$

has cardinality at most $N_D(2^{-j}, \Pi, \{X_i, \Gamma_i\})$. We also use the notation $A_0(\pi)(x) = 0$. For any index J , we clearly have

$$\pi = (\pi - A_J(\pi)) + \sum_{j=1}^J (A_j(\pi) - A_{j-1}(\pi)).$$

Notice that, for any π and j ,

$$\begin{aligned} \frac{1}{n} \text{Var} \left[\sum_{i=1}^n \Gamma_i Z_i (A_j(\pi)(X_i) - A_{j+1}(\pi)(X_i)) \mid \{X_i, \Gamma_i\}_{i=1}^n \right] \\ = 4\widehat{V} D^2 (A_j(\pi)(X_i), A_{j+1}(\pi)(X_i)) \leq 2^{2-2j} \widehat{V}, \end{aligned} \quad (36)$$

where $\widehat{V} = \sum_{i=1}^n \Gamma_i^2 / n$. Now, for every $n, j \geq 1$ and a sequence $\delta_n > 0$, define the event

$$\begin{aligned} \mathcal{E}_{j,n} := \left\{ \sup_{\pi \in \Pi} \left| \frac{1}{\sqrt{n}} \sum_{i=1}^n \Gamma_i Z_i (A_j(\pi)(X_i) - A_{j+1}(\pi)(X_i)) \right| \right. \\ \left. \geq 2^{2-j} \sqrt{\widehat{V}} \sqrt{\log(N_H(4^{-j}, \Pi)) + \log\left(\frac{2j^2}{\delta_n}\right)} \right\} \end{aligned} \quad (37)$$

Recall that, by Bernstein's inequality, for any independent, mean-zero variables S_i with $|S_i| \leq M$, and any constant $t > 0$

$$\mathbb{P} \left[\frac{1}{\sqrt{n}} \left| \sum_{i=1}^n S_i \right| \geq 2t \sqrt{\frac{1}{n} \sum_{i=1}^n \mathbb{E}[S_i^2]} \right] < 2e^{-t^2}, \text{ for all } 0 < t \leq \frac{1}{2M} \sqrt{\sum_{i=1}^n \mathbb{E}[S_i^2]}. \quad (38)$$

We can use this inequality to bound the probability of $\mathcal{E}_{j,n}$ if we condition on the X_i and Γ_i , use (36), apply a union bound, and note that $N_D(2^{-j}, \Pi, \{X_i, \Gamma_i\}) \leq N_H(4^{-j}, \Pi)$ thanks to (35):

$$\mathbb{P}[\mathcal{E}_{j,n}] \leq \frac{\delta_n}{j^2} + \mathbb{P} \left[\sqrt{\log(N_H(4^{-j}, \Pi)) + \log\left(\frac{2j^2}{\delta_n}\right)} \leq \frac{\sqrt{n\widehat{V}}}{2 \sup_{\{1 \leq i \leq n\}} \{|\Gamma_i|\}} \right]. \quad (39)$$

Here, the first right-hand-side term δ_n/j^2 is the failure probability in Bernstein's inequality, while the condition inside the $\mathbb{P}[\cdot]$ in the second summand is the precondition to applying Bernstein's inequality in (38).

At a high level, the inequality (39) shows that we can get bounds for the behavior of our stochastic process of interest for the first few terms in our chain of approximating sets. Recall that, by Assumption 2, we know that

$$\sqrt{\log(N_H(2^{-2j}, \Pi))} = \mathcal{O}\left(2^{(1-\omega)j}\right),$$

for some small $\omega > 0$. Thanks to this assumption, and the fact that

$$\mathbb{P} \left[\sqrt{n\widehat{V}} / \sup \{|\Gamma_i|\} \leq \sqrt{\frac{n}{\log(n)}} \right]$$

goes to 0 exponentially fast because Γ_i is sub-Gaussian, we can verify that

$$\mathbb{P} \left[\sqrt{\log(N_H(4^{-J(n)}, \Pi)) + \log\left(\frac{2j(n)^2}{\delta_n}\right)} \leq \frac{\sqrt{n\widehat{V}}}{2 \sup_{\{1 \leq i \leq n\}} \{|\Gamma_i|\}} \right], \quad (40)$$

also goes to zero exponentially fast for $J(n) := \lfloor \log_2(n)(1 + \omega)/2 \rfloor$, and any choice of δ_n that decays slower than $1/n$. Thus, with this choice of $J(n)$, the contribution of the second summand in (39) is negligible in comparison with the main term δ_n/j^2 .

We are now ready pull everything together. First, thanks to (39) and (40), we find that the event $\mathcal{E}_{j,n}$ holds for all $1 \leq j \leq J(n)$ with probability at least $1 - \delta_n \sum_{j=1}^{\infty} j^{-2} + o(1/n)$, and so

$$\begin{aligned} \sqrt{n} \sup_{\pi \in \Pi} \left| \frac{1}{n} \sum_{i=1}^n \Gamma_i Z_i \sum_{j=1}^{J(n)} (A_j(\pi) - A_{j-1}(\pi))(X_i) \right| \\ \leq 2\sqrt{\widehat{V}} \sum_{j=1}^{J(n)} 2^{2-j} \sqrt{\log(N_H(4^{-j}, \Pi)) + \log(j^2/\delta_n)} \\ \leq 8\sqrt{\widehat{V}} \int_0^1 \sqrt{\log(N_H(\varepsilon^2, \Pi)) + 2\log(\log(\varepsilon)) + \log(\delta_n^{-1})} d\varepsilon \\ \leq 8\sqrt{\widehat{V}} \left(\kappa(\Pi) + 2 \int_0^{e^{-1}} \sqrt{\log \log \varepsilon^{-1}} d\varepsilon + \sqrt{\log(\delta_n^{-1})} \right). \end{aligned} \quad (41)$$

Applying this bound separately for the sequences $\delta_n = \max\{2^{-k}, 1/n\}$ for $k = 1, 2, \dots$, we can turn the above into a bound on the expectation:

$$\begin{aligned} \sqrt{n} \mathbb{E} \left[\sup_{\pi \in \Pi} \left| \frac{1}{n} \sum_{i=1}^n \Gamma_i Z_i \sum_{j=1}^{J(n)} (A_j(\pi) - A_{j-1}(\pi))(X_i) \right| \right] \\ \leq 8\mathbb{E} \left[\sqrt{\widehat{V}} \right] \left(\kappa(\Pi) + 2 \int_0^{e^{-1}} \sqrt{\log \log \varepsilon^{-1}} d\varepsilon + \sum_{k=1}^{\infty} \frac{\sqrt{k \log(2)}}{2^k} \right) + \mathcal{O} \left(\frac{\log(n)}{\sqrt{n}} \right), \end{aligned}$$

where the last term is a crude bound on the expectation of our statistic of interest on the event that all Bernstein bounds fail, which occurs with probability at most $(1 + o(1))/n$. To bound this last term, we used the fact that Γ_i is sub-Gaussian, and so $\sup_{1 \leq i \leq n} \{|\Gamma_i|\}$ concentrates around a threshold scaling at most as $\sqrt{\log(n)}$. Finally, to obtain a bound on the expected Rademacher complexity, we still need to bound the discrepancy between our policy $\pi(X_i)$ and the finest approximation $A_{J(n)}(\pi)(X_i)$. We can use Jensen's inequality to check that

$$\begin{aligned} \frac{1}{n} \sum_{i=1}^n \Gamma_i Z_i (\pi(X_i) - A_{J(n)}(\pi)(X_i)) &\leq \sqrt{\frac{1}{n} \sum_{i=1}^n \Gamma_i^2 (\pi(X_i) - A_{J(n)}(\pi)(X_i))^2} \\ &= 2D(\pi(X_i), A_{J(n)}(\pi)(X_i)) \sqrt{\widehat{V}} \leq 2^{1-J(n)} \sqrt{\widehat{V}}, \end{aligned}$$

and so, given the above choice of $J(n) = \lfloor \log_2(n)(1 + \omega)/2 \rfloor$ for some $\omega > 0$,

$$\lim_{n \rightarrow \infty} \sqrt{n} \mathbb{E} \left[\frac{1}{n} \sum_{i=1}^n \Gamma_i Z_i (\pi(X_i) - A_{J(n)}(\pi)(X_i)) \right] = 0,$$

and this term is in fact negligible. Thus, we conclude that

$$\mathcal{R}(\Pi) \leq 8\sqrt{\frac{\mathbb{E}[\Gamma^2]}{n}} \left(\kappa(\Pi) + 2 \int_0^{e^{-1}} \sqrt{\log \log \varepsilon^{-1}} d\varepsilon + \sum_{k=1}^{\infty} \frac{\sqrt{k \log(2)}}{2^k} \right) + \mathcal{O}\left(\frac{\log(n)}{n}\right),$$

noting that $\mathbb{E}[\sqrt{\widehat{V}}] \leq \sqrt{\mathbb{E}[\Gamma^2]}$ by concavity of the square-root function.

5.2 Proof of Theorem 3

First, as argued by, e.g., [Bartlett and Mendelson \(2002\)](#) in the proof of their Theorem 8,

$$\mathbb{E} \left[\sup_{\pi \in \Pi} \left| \tilde{Q}(\pi) - Q(\pi) \right| \right] \leq 2\mathbb{E}[\mathcal{R}(\Pi)]. \quad (42)$$

Thus, to make use of Lemma 2, it suffices to bound $\sup_{\pi \in \Pi} \left| \tilde{Q}(\pi) - Q(\pi) \right|$ in terms of its expectation. Now, recall that $\tilde{Q}(\pi) = n^{-1} \sum \Gamma_i \pi(X_i)$, and that the Γ_i are sub-Gaussian. Now, because the Γ_i are not bounded, it is convenient to define truncated statistics

$$\tilde{Q}^{(-)}(\pi) = \frac{1}{n} \sum_{i=1}^n \Gamma_i^{(-)} \pi(X_i), \quad \Gamma_i^{(-)} = \Gamma_i \mathbf{1}(\{|\Gamma_i| \leq \log(n)\}).$$

Here, we of course have that $|\Gamma_i^{(-)}| \leq \log(n)$, and so we can apply Talagrand's inequality as described in [Bousquet \(2002\)](#) to these truncated statistics. We see that, for any $\delta > 0$, with probability at least $1 - \delta$,

$$\begin{aligned} \sup_{\pi \in \Pi} \left| \tilde{Q}^{(-)}(\pi) - Q^{(-)}(\pi) \right| &\leq \mathbb{E} \left[\sup_{\pi \in \Pi} \left| \tilde{Q}^{(-)}(\pi) - Q^{(-)}(\pi) \right| \right] \\ &+ \sqrt{2 \frac{\log(\delta)}{n} \left(\mathbb{E}[\Gamma_i^2] + 2 \log(n) \mathbb{E} \left[\sup_{\pi \in \Pi} \left| \tilde{Q}^{(-)}(\pi) - Q^{(-)}(\pi) \right| \right] \right)} + \frac{\log(n) \log(\delta)}{3n}, \end{aligned}$$

where we used the short-hand $Q^{(-)}(\pi) = \mathbb{E}[\tilde{Q}^{(-)}(\pi)]$. Moreover, because the Γ_i are sub-Gaussian, we can immediately verify that

$$\mathbb{E} \left[\left| \sup_{\pi \in \Pi} \left| \tilde{Q}^{(-)}(\pi) - Q^{(-)}(\pi) \right| - \sup_{\pi \in \Pi} \left| \tilde{Q}(\pi) - Q(\pi) \right| \right| \right]$$

decays exponentially fast in n . Using (42) and noting that, by Lemma 2, $\mathbb{E}[\mathcal{R}(\Pi)]$ decays as $\mathcal{O}(1/\sqrt{n})$, we conclude that with probability at least $1 - \delta$,

$$\sup_{\pi \in \Pi} \left| \tilde{Q}(\pi) - Q(\pi) \right| \leq 2\mathbb{E}[\mathcal{R}(\Pi)] + \sqrt{\frac{2\mathbb{E}[\Gamma_i^2] \log(\delta)}{n}} + \mathcal{O}\left(\frac{\log(n)}{n}\right), \quad (43)$$

thus establishing the first part of the theorem statement. Meanwhile, to prove the second part, we simply note that $\tilde{Q}(\tilde{\pi}) \geq \tilde{Q}(\pi^*)$ by construction, and then apply (43) at both $\tilde{\pi}$ and π^* .

5.3 Proof of Lemma 4

For any fixed policy π , we begin by expanding out the difference of interest. Write

$$\widehat{Q}_{+1}(\pi) = \frac{1}{n} \sum_{i=1}^n \pi(X_i) \left(\hat{\mu}_{+1}^{(-k(i))}(X_i) + \mathbf{1}(\{W_i = 1\}) \frac{Y_i - \hat{\mu}_{+1}^{(-k(i))}(X_i)}{\hat{e}_{+1}^{(-k(i))}(X_i)} \right),$$

and define $\widehat{Q}_{-1}(\pi)$ and $\widetilde{Q}_{\pm 1}(\pi)$ analogously, such that $\widehat{Q}(\pi) = \widehat{Q}_{+1}(\pi) - \widehat{Q}_{-1}(\pi)$, etc. Then,

$$\begin{aligned} \widehat{Q}_{+1}(\pi) - \widetilde{Q}_{+1}(\pi) &= \frac{1}{n} \sum_{i=1}^n \pi(X_i) \left(\hat{\mu}_{+1}^{(-k(i))}(X_i) - \mu_{+1}(X_i) \right. \\ &\quad \left. + \mathbf{1}(\{W_i = +1\}) \left(\frac{Y_i - \hat{\mu}_{+1}^{(-k(i))}(X_i)}{\hat{e}_{+1}^{(-k(i))}(X_i)} - \frac{Y_i - \mu_{+1}(X_i)}{e_{+1}(X_i)} \right) \right) \\ &= \frac{1}{n} \sum_{i=1}^n \pi(X_i) \left(\hat{\mu}_{+1}^{(-k(i))}(X_i) - \mu_{+1}(X_i) \right) \left(1 - \frac{\mathbf{1}(\{W_i = 1\})}{e_{+1}(X_i)} \right) \\ &\quad + \frac{1}{n} \sum_{\{i: W_i = 1\}} \pi(X_i) (Y_i - \mu_{+1}(X_i)) \left(\frac{1}{\hat{e}_{+1}^{(-k(i))}(X_i)} - \frac{1}{e_{+1}(X_i)} \right) \\ &\quad + \frac{1}{n} \sum_{\{i: W_i = 1\}} \pi(X_i) \left(\mu_{+1}(X_i) - \hat{\mu}_{+1}^{(-k(i))}(X_i) \right) \left(\frac{1}{\hat{e}_{+1}^{(-k(i))}(X_i)} - \frac{1}{e_{+1}(X_i)} \right). \end{aligned}$$

Denote these three summands by $A_{+1}(\pi)$, $B_{+1}(\pi)$, $C_{+1}(\pi)$. We will be to bound all 3 summands separately.

To bound the first term, it is helpful separate out the contributions of the K different folds:

$$A_{+1}^{(k)}(\pi) = \frac{1}{n} \sum_{\{i: k(i)=k\}} \pi(X_i) \left(\hat{\mu}_{+1}^{(-k)}(X_i) - \mu_{+1}(X_i) \right) \left(1 - \frac{\mathbf{1}(\{W_i = 1\})}{e_{+1}(X_i)} \right). \quad (44)$$

Now, because $\hat{\mu}_{+1}^{(-k)}(\cdot)$ was only computed using data from the $K-1$ folds, we can condition on the value of this function estimate to make the individual terms in the above sum independent. By Assumption 1, we know that

$$\sup_{x \in \mathcal{X}} \left| \left(\hat{\mu}_{+1}^{(-k)}(x) - \mu_{+1}(x) \right) \right| \leq \eta$$

with probability tending to 1, and so, by overlap, the individual summands in (44) are bounded by 1 with probability tending to 1. Thus, writing

$$V(k) = \mathbb{E} \left[\left(\hat{\mu}_{+1}^{(-k)}(X) - \mu_{+1}(X) \right)^2 \left(1 - \frac{\mathbf{1}(\{W_i = 1\})}{e_{+1}(X_i)} \right)^2 \mid \hat{\mu}_{+1}^{(-k)}(\cdot) \right],$$

we can can apply Theorem 3 to establish that

$$\frac{n}{n_k} \sup_{\pi \in \Pi} \left| A_{+1}^{(k)}(\pi) \right| \left| \hat{\mu}_{+1}^{(-k)}(\cdot) \right| = \mathcal{O}_P \left(\sqrt{\frac{V(k) \kappa(\Pi)}{n_k}} \right) + \mathcal{O} \left(\frac{\log(n_k)}{n_k} \right), \quad (45)$$

where $n_k = |\{i : k(i) = k\}|$ denotes the number of observations in the k -th fold. We constructed our double machine learning estimator using a finite number of evenly-sized folds, and so $n_k/n \rightarrow 1/K$. Thus, using Assumption 1 to bound the expected risk of $\hat{\mu}_{+1}^{(-k)}$, and applying (45) separately to all K folds, we conclude that

$$A_{+1}(\pi) = \mathcal{O}_P \left(n^{-\frac{3}{4}} \right).$$

We can use exactly the same argument to bound $B_{+1}(\pi)$ at the same rate.

It now remains to bound the final term, $C_{+1}(\pi)$. Here, we can use the Cauchy-Schwarz inequality to verify that

$$\begin{aligned} & \frac{1}{n} \sum_{\{i: W_i=1\}} \pi(X_i) \left(\mu_{+1}(X_i) - \hat{\mu}_{+1}^{-k(i)}(X_i) \right) \left(\frac{1}{\hat{e}_{+1}^{-k(i)}(X_i)} - \frac{1}{e_{+1}(X_i)} \right) \\ & \leq \sqrt{\frac{1}{n} \sum_{\{i: W_i=1\}} \left(\mu_{+1}(X_i) - \hat{\mu}_{+1}^{-k(i)}(X_i) \right)^2} \sqrt{\frac{1}{n} \sum_{\{i: W_i=1\}} \left(\frac{1}{\hat{e}_{+1}^{-k(i)}(X_i)} - \frac{1}{e_{+1}(X_i)} \right)^2}. \end{aligned}$$

Thus, given the L_2 -convergence rates of the machine learning methods from Assumption 1, as well as a second application of Cauchy-Schwarz, we find that

$$\mathbb{E} [C_{+1}(\pi)] \leq a \left(\left\lfloor \frac{K-1}{K} n \right\rfloor \right) / \sqrt{\left\lfloor \frac{K-1}{K} n \right\rfloor};$$

note that, here, we had to account for the fact each estimator was computed while omitting one of the K data folds. The desired conclusion follows from Markov's inequality, along with an application of the same argument to $\hat{Q}_{-1}(\pi)$.

5.4 Proof of Theorem 6

To establish this result, we follow the strategy in the proof of Theorem 3. We apply Talagrand's inequality to get the following analogue to (43),

$$\sup \left\{ \left| \tilde{Q}(\pi) - Q(\pi) \right| : \pi \in \Pi_\lambda \right\} \leq 2\mathbb{E} [\mathcal{R}(\Pi_\lambda)] + \sqrt{\frac{2V_\lambda \log(\delta)}{n}} + \mathcal{O} \left(\frac{\log(n)}{n} \right), \quad (46)$$

and so our task again reduces to bounding the Rademacher complexity $\mathbb{E} [\mathcal{R}(\Pi_\lambda)]$. At this point, however, it is convenient to slightly alter our definition of Rademacher complexity, and set

$$\mathcal{R}(\Pi_\lambda) := \sup \left\{ \frac{1}{n} \sum_{i=1}^n Z_i (\Gamma_i \pi(X_i) - Q(\pi^*)) \right\}, \quad (47)$$

where the Z_i are independent Rademacher variables and the Γ_i are defined as before. Here, the addition of a constant offset by no means alters the argument behind (46); formally, we would get to this notion of Rademacher complexity by trying to establish concentration of $|\tilde{Q}(\pi) - Q(\pi^*) - (Q(\pi) - Q(\pi^*))|$. However, this offset term will let us leverage the fact that $\pi \in \Pi_\lambda$ to get better bounds.

Now, we can start following the proof of Lemma 2. Recall that the first step of the proof involved constructing a sequence of approximating policies $A_j(\pi) : \mathcal{X} \rightarrow \{\pm 1\}$. These

approximating policies had the property that $D(A_j(\pi), A_{j+1}(\pi)) \leq 2^{-j}$ for all policies $\pi \in \Pi$, with D as defined in (34), and that the set $\{A_j(\pi) : \pi \in \Pi\}$ has cardinality at most $N_H(4^{-j}, \Pi)$.

Here, we need a slightly more strict construction. A simple extension of the same argument shows that we can create an analogous set of approximating policies $A_j^\lambda(\pi) \in \Pi_\lambda$, such that $D(A_j^\lambda(\pi), A_{j+1}^\lambda(\pi)) \leq 2^{-j}$ for all policies $\pi \in \Pi$, and the cardinality of the set of approximating policies is bounded by

$$|\{A_j^\lambda(\pi) : \pi \in \Pi_\lambda\}| \leq N_H(4^{-(j+1)}, \Pi_\lambda) \leq N_H(4^{-(j+1)}, \Pi).$$

Note that, in taking differences between policies, the constant offset $Q(\pi^*)$ in (47) gets canceled out; and so the argument can proceed verbatim this far. We can also establish concentration of the discrepancies between approximating policies $\sum_i Z_i \Gamma_i (A_j(\pi) - A_{j-1}(\pi))(X_i)$, and get the same result as in (39), while noting that we need to change the 4^{-j} to a $4^{-(j+1)}$ in the argument to N_H . Note that these discrepancy bounds still scale as $\sum_{i=1}^n \Gamma_i^2/n \sim V_{\max}$; and, moreover, these rates arise from the geometry of the problem (rather than as variance bounds), so there is no obvious way to replace these terms with something scaling as V_λ .

In order to leverage the fact that $\pi \in \Pi_\lambda$, we need to establish Rademacher concentration for the approximating policies themselves, rather than for differences between policies. To do so, we first note that for any policy $\pi \in \Pi_\lambda$

$$\text{Var} [Z_i (\Gamma_i \pi(X_i) - Q(\pi^*)) \mid \{X_i, \Gamma_i\}] = \frac{1}{n} \sum_{i=1}^n (\Gamma_i \pi(X_i) - Q(\pi^*))^2,$$

and, moreover, we can check that

$$\begin{aligned} \mathbb{E} [(\Gamma_i \pi(X_i) - Q(\pi^*))^2] &= \text{Var} [\Gamma_i \pi(X_i)] + (\mathbb{E} [\Gamma_i \pi(X_i)] - Q(\pi^*))^2 \\ &= V(\pi) + R(\pi)^2 = V_* + 2Q(\pi^*)R(\pi) \leq V_\lambda, \end{aligned}$$

where all above algebraic manipulations follow immediately from the definitions of the involved quantities (e.g., recall that $\bar{Q}(\pi) = \sum_i \pi(X_i) \Gamma_i/n$). Given these preliminaries, a straight-forward application of Bernstein's inequality (38) tells us that, for any $\delta > 0$,

$$\begin{aligned} \limsup_{n \rightarrow \infty} \mathbb{P} \left[\frac{1}{\sqrt{nV_\lambda}} \sup_{\{\pi \in \Pi_\lambda\}} \left\{ \left| \sum_{i=1}^n Z_i (\Gamma_i A_j^\lambda(\pi)(X_i) - Q(\pi^*)) \right| \right\} \right. \\ \left. \geq 2\sqrt{\log(N_H(4^{-(j+1)}, \Pi)) + \log\left(\frac{\delta}{2}\right)} \right] \leq \delta. \end{aligned}$$

Furthermore, continuing the same reasoning as before, we see that

$$\begin{aligned} \limsup_{n \rightarrow \infty} \frac{1}{\sqrt{nV_\lambda}} \mathbb{E} \left[\sup_{\{\pi \in \Pi_\lambda\}} \left\{ \left| \sum_{i=1}^n Z_i (\Gamma_i A_j^\lambda(\pi)(X_i) - Q(\pi^*)) \right| \right\} \right] \\ \leq 2\sqrt{\log(N_H(4^{-(j+1)}, \Pi))} + C, \\ \leq 2\sqrt{\alpha(j+1)\log(4)} + C, \end{aligned} \tag{48}$$

for some universal constant C , where on the last line we used our assumption (32).

It now only remains to plug our new bound (48) into the decomposition argument starting at (41), as follows. Given any fixed starting index J_0 , and defining $J(n)$ as in the proof of Lemma 2, we get

$$\begin{aligned} \sqrt{n} \sup_{\pi \in \Pi_\lambda} \left| \frac{1}{n} \sum_{i=1}^n Z_i (\Gamma_i \pi(X_i) - Q(\pi^*)) \right| &\leq \sqrt{n} \sup_{\pi \in \Pi_\lambda} \left| \frac{1}{n} \sum_{i=1}^n Z_i (\Gamma_i A_{J_0}^\lambda(\pi)(X_i) - Q(\pi^*)) \right| \\ &\quad + \sqrt{n} \sup_{\pi \in \Pi_\lambda} \left| \frac{1}{n} \sum_{i=1}^n Z_i \Gamma_i \sum_{j=J_0}^{J(n)} (A_j^\lambda(\pi)(X_i) - A_{j-1}^\lambda(\pi)(X_i)) \right| \\ &\quad + \sqrt{n} \sup_{\pi \in \Pi_\lambda} \left| \frac{1}{n} \sum_{i=1}^n Z_i \Gamma_i (\pi(X_i) - A_{J(n)}^\lambda(\pi)(X_i)) \right| \end{aligned}$$

We have already bounded the expectation of the first summand in (48), and, as argued in the proof of Lemma 2, the third summand is asymptotically negligible. Finally, following the of Lemma 2, for large enough n , the expectation of the second summand is bounded by

$$\begin{aligned} &8\sqrt{V_{\max}} \left(\int_0^{2^{1-J_0}} \sqrt{\log(N_H(4\varepsilon^2, \Pi))} + 2\sqrt{\log \log(\varepsilon^{-1})} + C \right) d\varepsilon \\ &\leq 8\sqrt{V_{\max}} \left(\int_0^{2^{1-J_0}} \sqrt{2\alpha \log(\varepsilon^{-1})} + 2\sqrt{\log \log(\varepsilon^{-1})} d\varepsilon + 2^{1-J_0} C \right), \end{aligned}$$

for some constant C ; to establish the above inequality, we also used (32). We can then verify by calculus that the above expression can be bounded by

$$C \sqrt{\max\{\alpha, 1\} V_{\max} 4^{1-J_0} \log(2^{1-J_0})}$$

for some (new) constant C . To show this, it is helpful to note that $\int_0^t \sqrt{\log(1/\varepsilon)} d\varepsilon \leq 2t\sqrt{\log(1/t)}$ for any $0 < t < 1/2$. Pulling everything together, we have thus found that, for any choice $J_0 \in \{1, 2, \dots\}$, there is a constant C for which

$$\sqrt{n} \mathbb{E} \left[\sup_{\pi \in \Pi_\lambda} \left| \frac{1}{n} \sum_{i=1}^n Z_i (\Gamma_i \pi(X_i) - Q(\pi^*)) \right| \right] \leq C \sqrt{\max\{\alpha, 1\} (V_\lambda J_0 + V_{\max} J_0 4^{-J_0})}.$$

Setting $J_0 := \lfloor \log_4(V_{\max}/V_\lambda) \rfloor$ establishes a bound on Rademacher complexity $\mathbb{E}[\mathcal{R}(\Pi_\lambda)]$ that leads to the desired result.

References

- A. Agarwal, D. Hsu, S. Kale, J. Langford, L. Li, and R. Schapire. Taming the monster: A fast and simple algorithm for contextual bandits. In *Proceedings of The 31st International Conference on Machine Learning*, pages 1638–1646, 2014.
- S. Athey, G. W. Imbens, and S. Wager. Approximate residual balancing: De-biased inference of average treatment effects in high dimensions. *arXiv preprint arXiv:1604.07125*, 2016a.
- S. Athey, J. Tibshirani, and S. Wager. Solving heterogeneous estimating equations with gradient forests. *arXiv preprint arXiv:1610.01271*, 2016b.

- P. Auer, N. Cesa-Bianchi, Y. Freund, and R. E. Schapire. The nonstochastic multiarmed bandit problem. *SIAM Journal on Computing*, 32(1):48–77, 2002.
- P. L. Bartlett and S. Mendelson. Rademacher and Gaussian complexities: Risk bounds and structural results. *Journal of Machine Learning Research*, 3:463–482, 2002.
- P. L. Bartlett and S. Mendelson. Empirical minimization. *Probability Theory and Related Fields*, 135(3):311–334, 2006.
- P. L. Bartlett, O. Bousquet, and S. Mendelson. Local Rademacher complexities. *Annals of Statistics*, pages 1497–1537, 2005.
- H. Bastani and M. Bayati. Online decision-making with high-dimensional covariates. 2015.
- A. Belloni, V. Chernozhukov, I. Fernández-Val, and C. Hansen. Program evaluation with high-dimensional data. *Econometrica*, 85(1):233–298, 2017.
- A. Beygelzimer and J. Langford. The offset tree for learning with partial labels. In *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 129–138. ACM, 2009.
- P. Bickel, C. Klaassen, Y. Ritov, and J. Wellner. *Efficient and Adaptive Estimation for Semiparametric Models*. Springer-Verlag, 1998.
- L. Bottou, J. Peters, J. Q. Candela, D. X. Charles, M. Chickering, E. Portugaly, D. Ray, P. Y. Simard, and E. Snelson. Counterfactual reasoning and learning systems: The example of computational advertising. *Journal of Machine Learning Research*, 14(1):3207–3260, 2013.
- S. Boucheron, G. Lugosi, and P. Massart. *Concentration inequalities: A nonasymptotic theory of independence*. Oxford university press, 2013.
- O. Bousquet. A Bennett concentration inequality and its application to suprema of empirical processes. *Comptes Rendus Mathématique*, 334(6):495–500, 2002.
- L. Breiman, J. Friedman, C. J. Stone, and R. A. Olshen. *Classification and Regression Trees*. CRC press, 1984.
- G. Chen, D. Zeng, and M. R. Kosorok. Personalized dose finding using outcome weighted learning. *Journal of the American Statistical Association*, (just-accepted), 2016.
- V. Chernozhukov, D. Chetverikov, M. Demirer, E. Duffo, C. Hansen, and W. Newey. Double machine learning for treatment and causal parameters. *arXiv preprint arXiv:1608.00060*, 2016.
- C. Cortes and V. Vapnik. Support-vector networks. *Machine learning*, 20(3):273–297, 1995.
- C. Cortes, Y. Mansour, and M. Mohri. Learning bounds for importance weighting. In *Advances in Neural Information Processing Systems*, pages 442–450, 2010.
- M. Dudík, J. Langford, and L. Li. Doubly robust policy evaluation and learning. In *Proceedings of the 28th International Conference on Machine Learning*, pages 1097–1104, 2011.
- R. M. Dudley. The sizes of compact subsets of Hilbert space and continuity of Gaussian processes. *Journal of Functional Analysis*, 1(3):290–330, 1967.
- E. Giné and V. Koltchinskii. Concentration inequalities and asymptotic results for ratio type empirical processes. *The Annals of Probability*, 34(3):1143–1216, 2006.
- B. S. Graham, C. C. D. X. Pinto, and D. Egel. Inverse probability tilting for moment condition models with missing data. *The Review of Economic Studies*, 79(3):1053–1079, 2012.
- T. Grubinger, A. Zeileis, and K.-P. Pfeiffer. evtrees: Evolutionary learning of globally optimal classification and regression trees in r. 61(1), 2014.

- J. Hahn. On the role of the propensity score in efficient semiparametric estimation of average treatment effects. *Econometrica*, pages 315–331, 1998.
- D. Haussler. Sphere packing numbers for subsets of the Boolean n-cube with bounded Vapnik-Chervonenkis dimension. *Journal of Combinatorial Theory, Series A*, 69(2):217–232, 1995.
- K. Hirano, G. W. Imbens, and G. Ridder. Efficient estimation of average treatment effects using the estimated propensity score. *Econometrica*, 71(4):1161–1189, 2003.
- K. Imai and M. Ratkovic. Estimating treatment effect heterogeneity in randomized program evaluation. *The Annals of Applied Statistics*, 7(1):443–470, 2013.
- N. Kallus. Learning to personalize from observational data. *arXiv preprint arXiv:1608.08925*, 2016.
- T. Kitagawa and A. Tetenov. Who should be treated? Empirical welfare maximization methods for treatment choice. Technical report, Centre for Microdata Methods and Practice, Institute for Fiscal Studies, 2015.
- A. Maurer and M. Pontil. Empirical Bernstein bounds and sample variance penalization. In *Conference on Learning Theory*, 2009.
- W. K. Newey. The asymptotic variance of semiparametric estimators. *Econometrica: Journal of the Econometric Society*, 62(6):1349–1382, 1994.
- J. Neyman. Sur les applications de la théorie des probabilités aux expériences agricoles: Essai des principes. *Roczniki Nauk Rolniczych*, 10:1–51, 1923.
- M. Qian and S. A. Murphy. Performance guarantees for individualized treatment rules. *Annals of Statistics*, 39(2):1180, 2011.
- J. Robins and A. Rotnitzky. Semiparametric efficiency in multivariate regression models with missing data. *Journal of the American Statistical Association*, 90(1):122–129, 1995.
- J. Robins, A. Rotnitzky, and L. Zhao. Analysis of semiparametric regression models for repeated outcomes in the presence of missing data. *Journal of the American Statistical Association*, 90(1):106–121, 1995.
- P. R. Rosenbaum and D. B. Rubin. The central role of the propensity score in observational studies for causal effects. *Biometrika*, 70(1):41–55, 1983.
- D. B. Rubin. Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of Educational Psychology*, 66(5):688, 1974.
- A. Schick. On asymptotically efficient estimation in semiparametric models. *The Annals of Statistics*, pages 1139–1151, 1986.
- A. Swaminathan and T. Joachims. Batch learning from logged bandit feedback through counterfactual risk minimization. *Journal of Machine Learning Research*, 16:1731–1755, 2015.
- R. Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, pages 267–288, 1996.
- M. J. van der Laan and S. Gruber. Collaborative double robust targeted maximum likelihood estimation. *The International Journal of Biostatistics*, 6(1), 2010.
- M. J. van der Laan and S. Rose. *Targeted Learning: Causal Inference for Observational and Experimental Data*. Springer Science & Business Media, 2011.
- V. Vapnik and A. Y. Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. *Theory of Probability and its Applications*, 16(2):264, 1971.
- S. Wager, W. Du, J. Taylor, and R. J. Tibshirani. High-dimensional regression adjustments in randomized experiments. *Proceedings of the National Academy of Sciences*, 113(45):12673–12678, 2016.

- B. Zhang, A. A. Tsiatis, M. Davidian, M. Zhang, and E. Laber. Estimating optimal treatment regimes from a classification perspective. *Stat*, 1(1):103–114, 2012.
- Y. Zhao, D. Zeng, A. J. Rush, and M. R. Kosorok. Estimating individualized treatment rules using outcome weighted learning. *Journal of the American Statistical Association*, 107(499):1106–1118, 2012.
- X. Zhou, N. Mayer-Hamblett, U. Khan, and M. R. Kosorok. Residual weighted learning for estimating individualized treatment rules. *Journal of the American Statistical Association*, (just-accepted):00–00, 2015.